Exhibit 30

CHART FOR U.S. PATENT NO. 8,417,871 ("the '871 Patent")

Accused Products:

Dell's products, including but not limited to Dell's PowerMax (*e.g.*, PowerMax 2000, 2500, 8000, and 8500), VMAX All Flash (*e.g.*, VMAX 250F, 450F, 850F, and 950F), and EMC VMAX (*e.g.*, VMAX 100K, 200K, and 400K) ("Accused Products"), infringe at least Claim 13 of the '871 Patent.

Claims	Exemplary Evidence of Infringement
13. An apparatus, comprising:	For example, the Accused Products comprise "modular storage components for compute and storage media" where the "compute modules are packaged as node pairs . . . [and e]ach node pair contains two PowerMax compute nodes, complete software and licensing." *See, e.g.*: 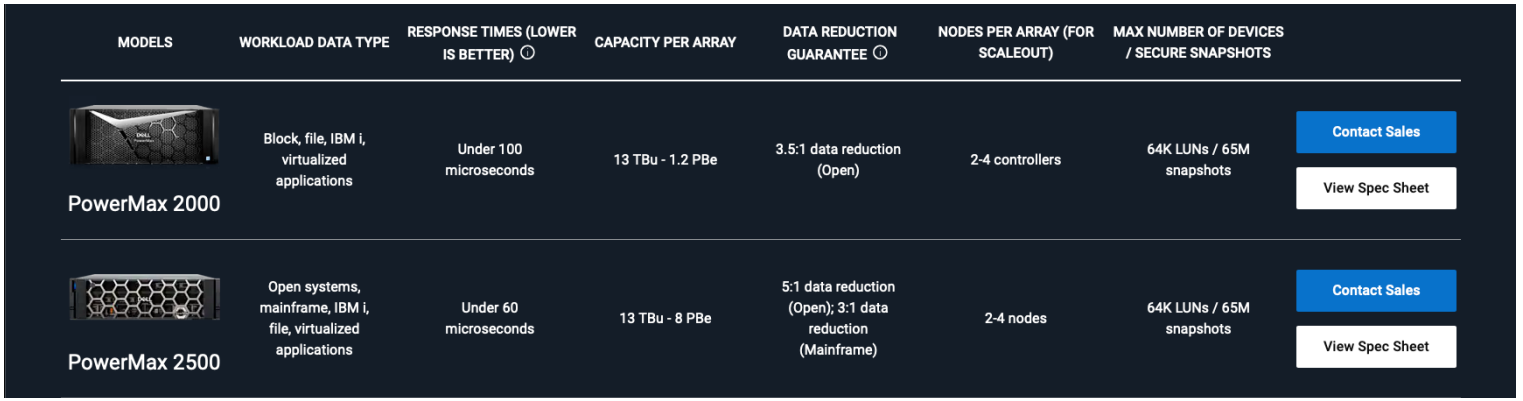

Claims	Exemplary Evidence of Infringement

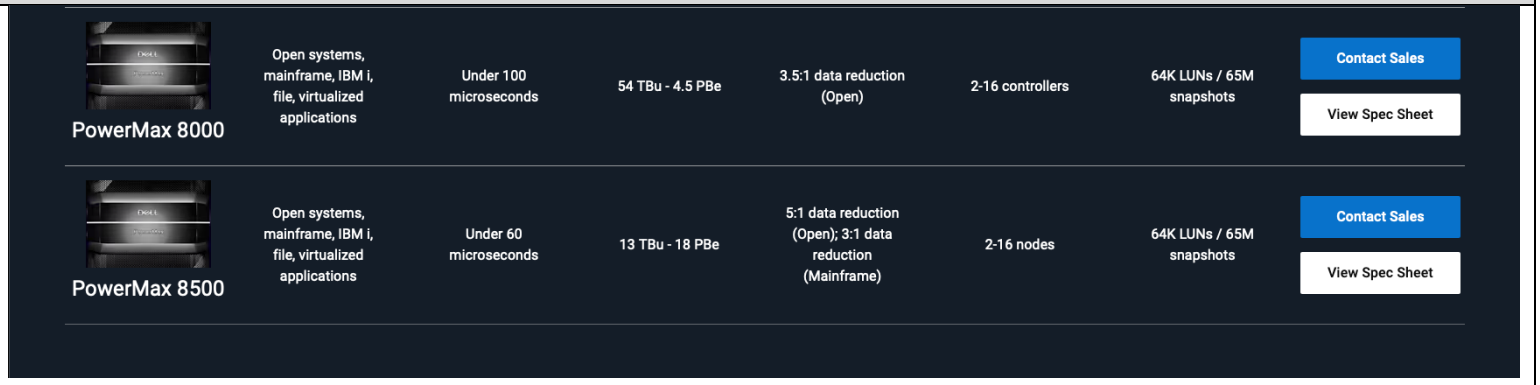

Servers, Storage & Networking – Storage – PowerMax

PowerMax is built from modular storage components for compute and storage media. The compute modules are packaged as node pairs. Each node pair contains two PowerMax compute nodes, complete software and licensing, cache memory, redundant power, and connectivity modules. These are combined with 48-slot Dynamic Media Enclosures (DMEs) to configure NVMe flash drives. PowerMax arrays are delivered with the Inclusive Software package. NVMe drive capacity can be added to the system to scale up to a total effective capacity of 8 PBe on the PowerMax 2500 and up to 18 PBe on the PowerMax 8500.

Specification Sheet – Dell PowerMax Array – DELL POWERMAX – Dell PowerMax 2500 and 8500

Claims	Exemplary Evidence of Infringement
	Data Storage – VMAX All Flash  Storage – EMC VMAX
13 [a] storage elements; and	The Accused Products comprise storage elements. For example, the Accused Products combine "all physical storage capacity . . . into Storage Resource Pools (SRPs)," which "consist of Disk Groups which contain [] a collection of hard drives." For example, "hard drives in each disk group are split into back-end data segments called TDATs" and "TDATs are placed into an associated Storage Tier." *See, e.g.*:

Claims	Exemplary Evidence of Infringement
	PowerMax Storage Resource Pools overview In PowerMax, all physical storage capacity is combined into Storage Resource Pools (SRPs). At the lowest levels, SRPs consist of Disk Groups which contain of a collection of hard drives sharing the same technology and performance characteristics. The hard drives in each disk group are split into individual back-end data device segments called TDATs. The TDATs are placed into an associated Storage Tier. An SRP is the collection of the total capacity of all its Storage Tiers – regardless of the underlying disk technology which the storage tiers are associated with. This physical capacity stored within an SRP is referred to its usable capacity (TBu). This usable capacity is accessed by hosts using thinly provisioned front-end storage devices called TDEVs. TDEVs are virtual representation of the SRP physical capacity which also considers overprovisioning and data reduction efficiencies. For example, an array with a single SRP which has 26 TBu, could be provisioned for 78 TB of host facing TDEV capacity when a data reduction ratio of 3:1 is applied. This 78 TB of virtualized host facing TDEV capacity is referred to be the effective capacity (TBe) of the SRP. When a PowerMax is sized, both the usable capacity and effective capacity are considered. The total usable capacity (TBu) is the primary driver for sizing hard-drive-layout configurations. The effective capacity (TBe) is a primary driver when sizing PowerMax cache. Host provisioned TDEVs to are placed into a storage group and assigned a Service Level. When a host writes application data to its provisioned TDEVs, this data is distributed across all the storage tiers within the SRP. Which storage tier the data is placed on within the SRP is governed by the Automated Data Placement (ADP) utility. ADP uses the PowerMax internal machine learning engine to employ predictive analytics and pattern recognition algorithms to place the data at the optimal physical location to ensure that the response time requirements for the assigned service level are met. Dell EMC PowerMax: Family Overview

Claims	Exemplary Evidence of Infringement
13 [b] a storage access system configured to:	The Accused Products comprise a storage access system. For example, the Accused Products include "Intel Xeon" "CPU[s]," "PowerMaxOS"/"HYPERMAX OS"/"ENGINUITY OS," and "are delivered with the Inclusive Software package." *See, e.g.*: 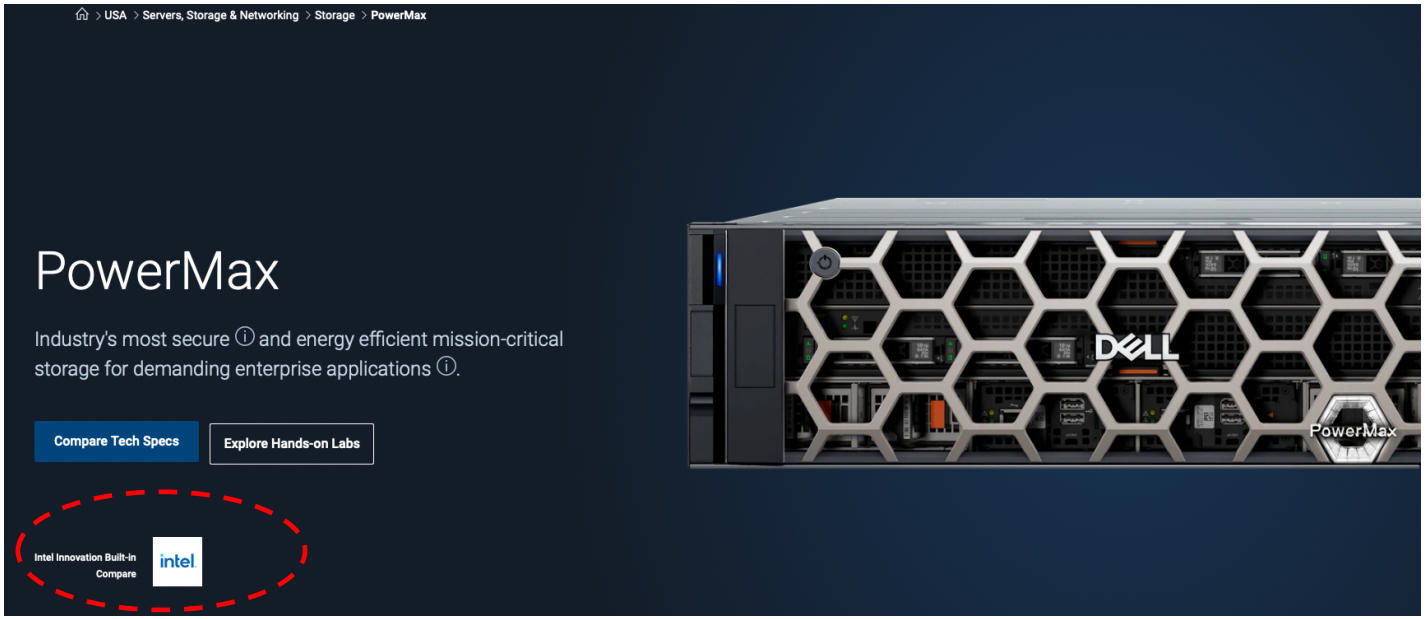 Servers, Storage & Networking – Storage – PowerMax

Claims	Exemplary Evidence of Infringement

Array family	PowerMax 2500	PowerMax 8500
Node Pairs		
NUMBER OF NODE PAIRS	1 to 2	1 to 8
NODE PAIR MODULE	3U	3U
CPU	Memory config 1-3: Intel Xeon Gold 5218 2.8 GHz with 16 core[1] Memory config 4: Intel Xeon Gold 6240L	Memory config 2-3: Intel Xeon Gold 6254 3.9 GHz with 18 core[1] Memory Config 4: Intel Xeon Gold 8280L
CORE NUMBER PER CPU/PER NODE PAIR/PER SYSTEM	Memcfg 1-3: 16/64/128 Memcfg 4: 18/72/144[5]	Memcfg 1-3: 18/72/576 Memcfg 4: 20/80/608[4,5]

Specification Sheet – Dell PowerMax Array – DELL POWERMAX – Dell PowerMax 2500 and 8500

Array family	PowerMax 2000	PowerMax 8000
Bricks/zBricks		
Number of Bricks or zBricks[5]	1 to 2	1 to 8
ENGINE ENCLOSURE	4u	4u
CPU	Intel Xeon E5-2650-v4 2.5 GHz 12 core[4]	Intel Xeon E5-2697-v4 2.8 GHz 18 core[4]
# CORES PER CPU/PER ENGINE/PER SYSTEM	12/48/96	18/72/576

Specification Sheet – POWERMAX FAMILY – PowerMax 2000 and 8000

Claims	Exemplary Evidence of Infringement

Data Storage – VMAX All Flash

Array family	VMAX 250F/VMAX 250FX	VMAX 950F/VMAX 950FX
V-BRICKS		
Number of V-Bricks	1 to 2	1 to 8
ENGINE ENCLOSURE	4u	4u
CPU	Intel Xeon E5-2650-v4 [4]2.5 GHz 12 core	Intel Xeon E5-2697-v4 [4]2.8 GHz 18 core
# CORES PER CPU/PER ENGINE/PER SYSTEM	12/48/96	18/72/576

Specification Sheet – VMAX ALL FLASH FAMILY – VMAX 250F, 950F

ARRAY FAMILY	VMAX 450F/VMAX 450FX	VMAX 850F/VMAX 850FX
V-BRICKS		
NUMBER OF V-BRICKS	1 to 4	1 to 8
ENGINE ENCLOSURE	4u	4u
CPU	Intel Xeon E5-2650-v2 [4]3.0 GHz 8 core	Intel Xeon E5-2697-v2 [4]3.0 GHz 12 core
# CORES PER CPU/PER ENGINE/PER SYSTEM	8/32/128	12/48/384

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>Specification Sheet – VMAX ALL FLASH FAMILY – VMAX 450F, 850F

VMAX3 FAMILY SPECIFICATIONS
<table>
<tr><th>VMAX3 ARRAY</th><th>VMAX 100K</th><th>VMAX 200K</th><th>VMAX 400K</th></tr>
<tr><td colspan="4">ENGINE</td></tr>
<tr><td>Number of Engines supported</td><td>1 to 2</td><td>1 to 4</td><td>1 to 8</td></tr>
<tr><td>Engine Enclosure</td><td>4u</td><td>4u</td><td>4u</td></tr>
<tr><td>CPU</td><td>Intel Xeon E5-2620-v2
2.1 GHz 6 core</td><td>Intel Xeon E5-2650-v2
2.6 GHz 8 core</td><td>Intel Xeon E5-2697-v2
2.7 GHz 12 core</td></tr>
<tr><td>Dynamic Virtual Matrix BW</td><td>700GB/s</td><td>700GB/s</td><td>1400GB/s</td></tr>
<tr><td># Cores per CPU/per Engine/per System</td><td>6/24/48</td><td>8/32/128</td><td>12/48/384</td></tr>
</table>
Specification Sheet – DELL EMC VMAX3 FAMILY – VMAX 100K, 200K, 400K

PowerMax is built from modular storage components for compute and storage media. The compute modules are packaged as node pairs. Each node pair contains two PowerMax compute nodes, complete software and licensing, cache memory, redundant power, and connectivity modules. These are combined with 48-slot Dynamic Media Enclosures (DMEs) to configure NVMe flash drives. PowerMax arrays are delivered with the Inclusive Software package. NVMe drive capacity can be added to the system to scale up to a total effective capacity of 8 PBe on the PowerMax 2500 and up to 18 PBe on the PowerMax 8500.

Specification Sheet – Dell PowerMax Array – DELL POWERMAX – Dell PowerMax 2500 and 8500</td></tr>
<tr><td>13 [b][i] perform write operations configured to write same data into the</td><td>The Accused Products perform write operations configured to write same data into the storage elements sequentially one at a time so a number of the storage elements remain available for read operations while the other storage elements are being written with the data, wherein the number of storage elements available for the read operations is associated with a selectable performance index.

For example, the Accused Products provide “PowerMax Data Services” including “[r]emote replication” with the “Symmetrix Remote Data Facility (SRDF)” and “[h]igh availability with SRDF/Metro.” For example, a “SRDF device is a logical device paired with another logical device that resides in a second array” and “[d]uring normal</td></tr>
</table>

Claims	Exemplary Evidence of Infringement
storage elements sequentially one at a time so a number of the storage elements remain available for read operations while the other storage elements are being written with the data, wherein the number of storage elements available for the read operations is associated with a selectable performance index;	operations, host I/O writes to the R1 device are mirrored over the SRDF links to the R2 device." For example, in "SRDF synchronous mode, an R2 device is typically in read-only mode (write disabled) that allows a remote host to read from the R2 devices" and a "host connected to the R2 device has read-only access to the R2 device." For example, the Accused Products "are preconfigured with . . . service levels" that are used to "specify the performance objectives" and "ensure that applications have consistent and predictable performance," where a "service level is the response time target for the storage group." For example, the "[t]arget response time" is the "average response time expected for the storage group based on the selected service level," and along with a "target response time, service levels also have either an upper response time limit or both an upper and lower response time limit." For example, "[u]sers may set the required service level objective (SLO) independently on both source and target SRDF/Metro paired arrays." *See, e.g.*: PowerMax Data Services help protect, manage, and move customer data on the array. These services run natively or embedded inside the PowerMax itself using the PowerMaxOS hypervisor to provide a resource abstraction layer. This allows the data services to share array resources — CPU cores, cache, and bandwidth. Doing this optimizes performance across the entire system and reduces complexity in the environment as resources do not need to be dedicated. Some of the most sought-after data services that are offered with the PowerMax product line are: • Remote replication with SRDF • High Availability with SRDF/Metro • Local replication with TimeFinder SnapVX • Embedded NAS (eNAS) • Embedded Unisphere for PowerMax (eManagement) Next-Generation PowerMax Family Overview

Claims	Exemplary Evidence of Infringement
	PowerMax Storage Resource Pools overview In PowerMax, all physical storage capacity is combined into Storage Resource Pools (SRPs). At the lowest levels, SRPs consist of Disk Groups which contain of a collection of hard drives sharing the same technology and performance characteristics. The hard drives in each disk group are split into individual back-end data device segments called TDATs. The TDATs are placed into an associated Storage Tier. An SRP is the collection of the total capacity of all its Storage Tiers – regardless of the underlying disk technology which the storage tiers are associated with. This physical capacity stored within an SRP is referred to its usable capacity (TBu). This usable capacity is accessed by hosts using thinly provisioned front-end storage devices called TDEVs. TDEVs are virtual representation of the SRP physical capacity which also considers overprovisioning and data reduction efficiencies. For example, an array with a single SRP which has 26 TBu, could be provisioned for 78 TB of host facing TDEV capacity when a data reduction ratio of 3:1 is applied. This 78 TB of virtualized host facing TDEV capacity is referred to be the effective capacity (TBe) of the SRP. When a PowerMax is sized, both the usable capacity and effective capacity are considered. The total usable capacity (TBu) is the primary driver for sizing hard-drive-layout configurations. The effective capacity (TBe) is a primary driver when sizing PowerMax cache. Host provisioned TDEVs to are placed into a storage group and assigned a Service Level. When a host writes application data to its provisioned TDEVs, this data is distributed across all the storage tiers within the SRP. Which storage tier the data is placed on within the SRP is governed by the Automated Data Placement (ADP) utility. ADP uses the PowerMax internal machine learning engine to employ predictive analytics and pattern recognition algorithms to place the data at the optimal physical location to ensure that the response time requirements for the assigned service level are met. Dell EMC PowerMax: Family Overview

Claims	Exemplary Evidence of Infringement
	The Symmetrix Remote Data Facility (SRDF) maintains real-time (or near real-time) copies of data on a production storage array at one or more remote storage arrays. SRDF has three primary applications: • Disaster recovery • High availability • Data migration Dell EMC SRDF Introduction

Claims	Exemplary Evidence of Infringement

An SRDF device is a logical device paired with another logical device that resides in a second array. The arrays are connected by SRDF links. R1 devices are the member of the device pair at the primary (production) site. R1 devices are generally read/write accessible to the host. R2 devices are the members of the device pair at the secondary (remote) site. During normal operations, host I/O writes to the R1 device are mirrored over the SRDF links to the R2 device.

Figure 1. Traditional SRDF device pair states

Traditionally, data on R2 devices are not available to the host while the SRDF relationship is active. In SRDF synchronous mode, an R2 device is typically in read-only mode (write disabled) that allows a remote host to read from the R2 devices. In a typical open systems host environment, the production host has read/write access to the R1 device. A host connected to the R2 device has read-only access to the R2 device. To access the R2 device of a traditional synchronous relationship, a manual failover or swap operation must be performed to write enable the R2 site to accept host writes.

Claims	Exemplary Evidence of Infringement
	Performance statistic exchange begins once the SRDF/Metro Active mode and ActiveActive or ActiveBias pair state have been achieved. Each side then incorporates the FAST statistics from the other side to ensure each side represents the workload as a whole (R1+R2 workload). Users may set the required service level objective (SLO) independently on both source and target SRDF/Metro paired arrays. There are currently no restrictions in this area as FAST data movement is transparent from SRDF/Metro. Dell PowerMax and VMAX All Flash: SRDF/Metro Overview and Best Practices A service level is the response time target for the storage group. The service level enables you set the storage array with the desired response time target for the storage group. It automatically monitors and adapts to the workload to maintain (or meet) the response time target. The service level includes an optional workload type. The optional workload type can be used to further tune expectations for the workload storage group to provide enough flash to meet your performance objective.

Claims	Exemplary Evidence of Infringement
	By default, storage systems running HYPERMAX OS 5977 or higher are preconfigured with a single Storage Resource Pool (SRP). The SRP contains all the hard disks on the system that is organized into disk groups by technology, capacity, rotational speed, and RAID protection type. Storage administrators can view all SRPs configured on the system, and the demand that storage groups are placing on them. Storage systems are also preconfigured with several service levels and workloads. Storage administrators use the service levels and workloads to specify the performance objectives for the application they are provisioning. When provisioning storage for an application, storage administrators assign the appropriate SRP, service level, and workload to the storage group containing the LUNs associated with the application. 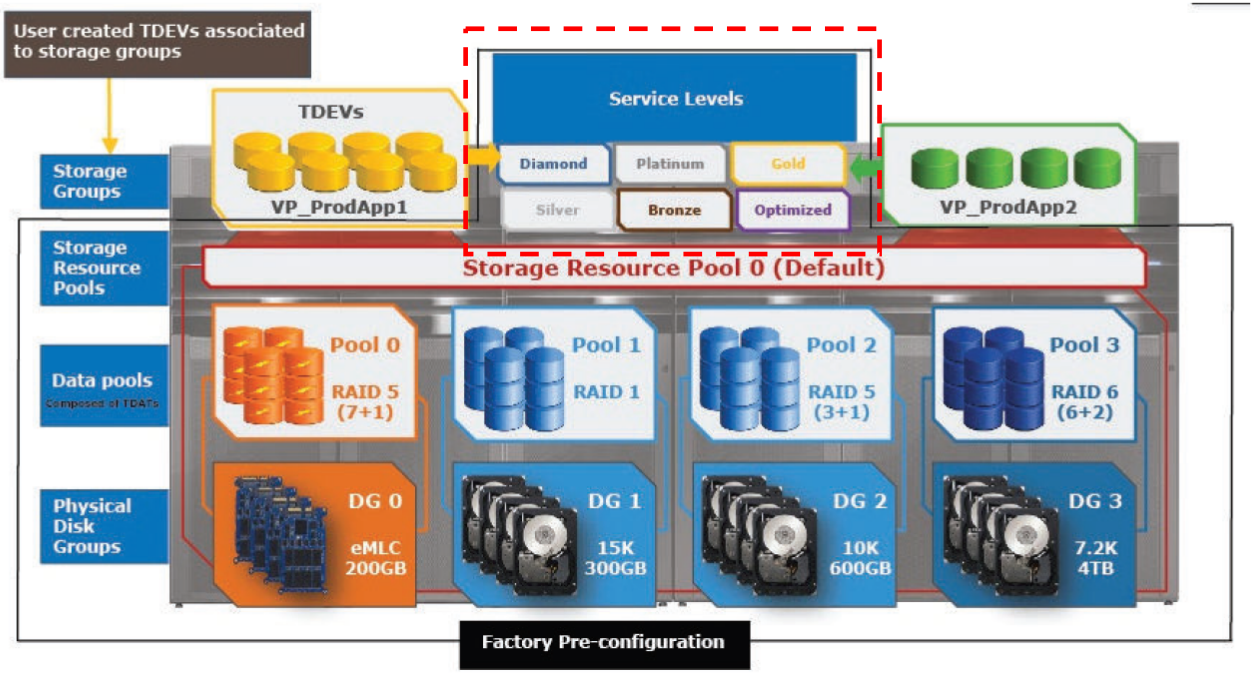

Claims	Exemplary Evidence of Infringement
	Dell EMC Unisphere for PowerMax Product Guide **PowerMaxOS:** Beginning with 5978, the operating environment run on PowerMax and VMAX All Flash systems. **Storage group (SG):** A logical grouping of thin devices that are provisioned and associated with a particular application. **Response time (RT):** The total amount of time it takes to respond to a request for service. **Target response time:** The average response time expected for the storage group based on the selected service level. **Upper response time limit:** The maximum response time specified by the selected service level. **Lower response time limit:** The minimum response time specified by the selected service level. Service levels for PowerMaxOS ensure that applications have consistent and predictable performance by allowing users to separate storage groups based on performance requirements and business importance. PowerMaxOS allows you to set specific service levels to ensure the highest-priority application response times are not impacted by lower-priority applications. The available service levels are defined in PowerMaxOS and can be applied at the creation of a storage group or can be modified to an existing storage group at any time.

Claims	Exemplary Evidence of Infringement

Service level options

Service levels are offered with various ranges of performance expectations which are defined by their own characteristics of a target response time. The target response time is the average response time expected for the storage group based on the selected service level. Along with a target response time, service levels also have either an upper response time limit or both an upper and lower response time limit.

The service levels offered are detailed in Table 1. All service levels shown, with the exception of Optimized, have a target response time.

Table 1 Service levels for PowerMaxOS

Service level	Diamond	Platinum	Gold	Silver	Bronze	Optimized
Target response time	✓	✓	✓	✓	✓	
Upper response time limit	✓	✓	✓	✓	✓	
Lower response time limit				✓	✓	

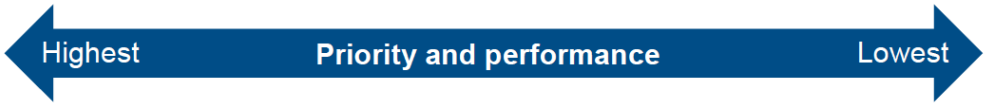

Diamond, Platinum, and Gold: These service levels have the highest priority and performance. Each has an upper response time limit but no lower response time limit which ensures they will be serviced as fast as possible.

Claims	Exemplary Evidence of Infringement

Silver and Bronze: These service levels have both an upper and lower limit designed to allow higher-priority service levels to be unaffected. These are managed such that their average response time will be greater than or equal to the lower response time limit.

Optimized: This service level does not have a target response time nor an upper or lower limit. Optimized is designed to use all allowable resources, equal to that of Diamond, and is not managed to assist any other service level. Storage groups set with any other service level will also not be managed to assist storage groups set Optimized. Optimized should be used on systems where application performance and consistency are not of relative importance and should not be mixed on systems with other service levels.

The following table shows the target response time for each service level along with an indication of the lower response time.

Table 2 Service level response times

Service level	Target response time	Lower response time
Diamond	0.6 ms*	None
Platinum	0.8 ms*	None
Gold	1 ms	None
Silver	3.6 ms	~3.6 ms
Bronze	7.2 ms	~7.2 ms
Optimized	N/A	N/A

Setting service levels

Service levels can be applied to a storage group when either creating a new storage group or by modifying an existing storage group. Users also have the ability to change service levels at any time to apply the desired response time performance expectation.

Claims	Exemplary Evidence of Infringement

Service level functionality

Service levels for PowerMaxOS are designed to put relative importance on application storage groups. This function allows users to manage applications based on predictability and consistency in response time. With the ability to manage the performance of the array, users can determine whether to utilize as much of the resources the array allows or as little of the resources regardless of the system capabilities. With this ability to manage performance, users assign importance on applications based on specific needs.

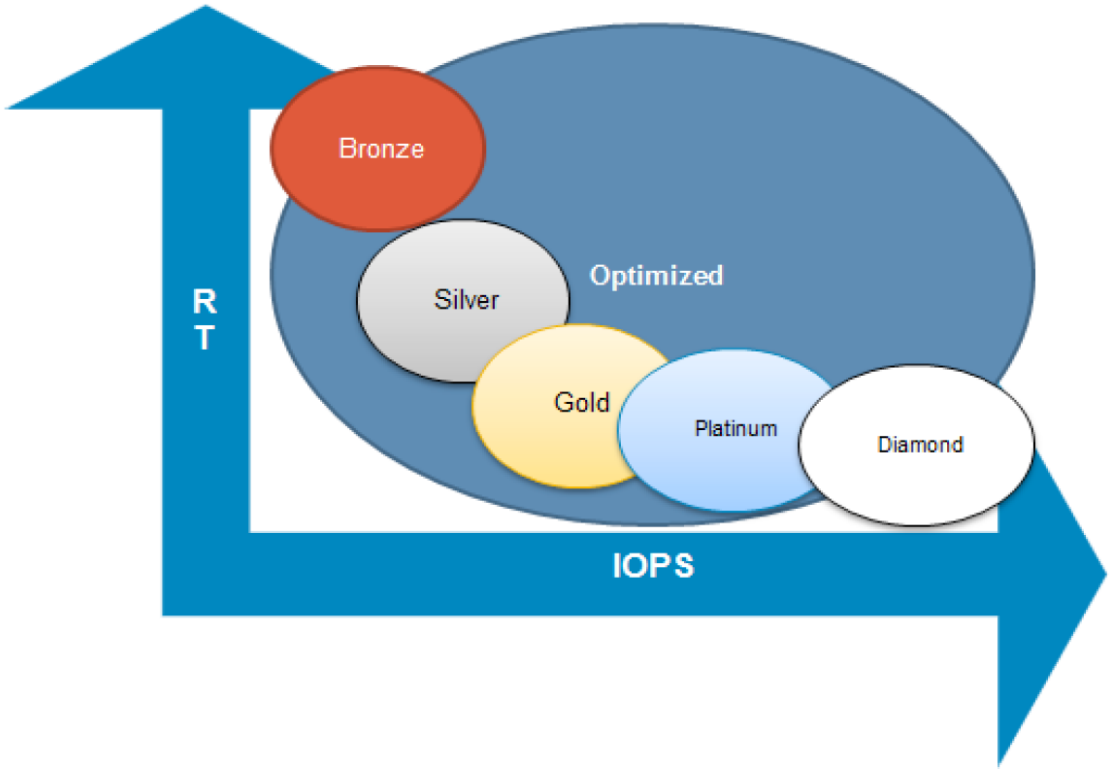

Figure 10 Relative response time compared to IOPS

Claims	Exemplary Evidence of Infringement
	Figure 13 illustrates service levels as they apply to the response-time upper and lower limits and shows how host I/O limits relate to setting the maximum allowable throughput IOPS. 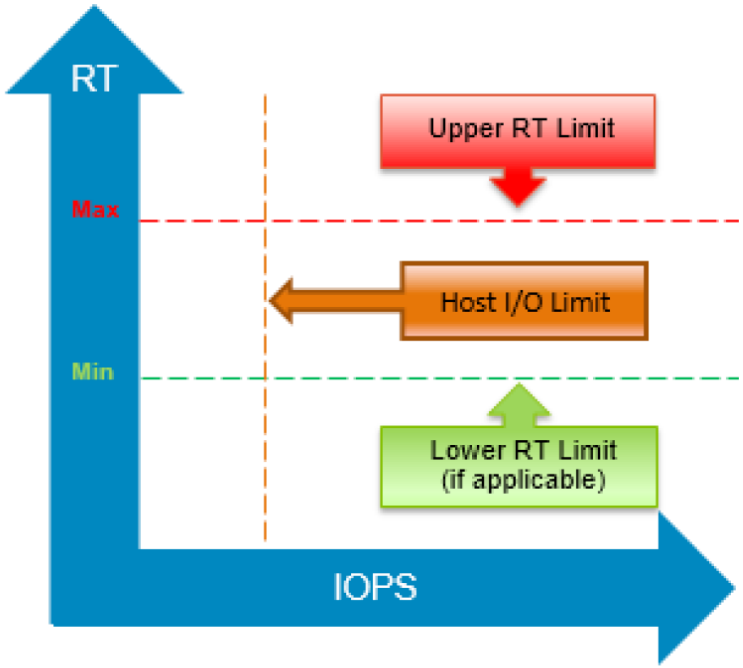 Figure 13 Service level management relative to host I/O limits Technical White Paper – Dell EMC PowerMax: Service Levels for PowerMaxOS
13 [b][ii] map read addresses for the read operations to multiple	As configured, the Accused Products map read addresses for the read operations to multiple different ones of the storage elements not currently being used for the write operations. For example, the Accused Products combine "all physical storage capacity into Storage Resource Pools (SRPs)," which "consist of Disk Groups which contain [] a collection of hard drives." For example, the "hard drives in each disk group are split into back-end data segments called TDATs" and "TDATs are placed into an associated Storage

Claims	Exemplary Evidence of Infringement
different ones of the storage elements not currently being used for the write operations; and	Tier." For example, the "useable capacity of" an "SRP" is "accessed by hosts using thinly-provisioned front-end devices called TDEVs," where "TDEVs are [a] virtual representation of the SRP physical capacity" and "are placed into a storage group and assigned a Service Level." For example, the Accused Products provide "PowerMax Data Services" including "[r]emote replication" with the "Symmetrix Remote Data Facility (SRDF)" and "[h]igh availability with SDRF/Metro." For example, a "SRDF device is a logical device paired with another logical device that resides in a second array" and "[d]uring normal operations, host I/O writes to the R1 device are mirrored over the SRDF links to the R2 device." For example, in "SRDF synchronous mode, an R2 device is typically in read-only mode (write disabled) that allows a remote host to read from the R2 devices" and a "host connected to the R2 device has read-only access to the R2 device." *See, e.g.*: • High Availability with SRDF/Metro • Local replication with TimeFinder SnapVX • Embedded NAS (eNAS) • Embedded Unisphere for PowerMax (eManagement) • Remote replication with SRDF • High Availability with SRDF/Metro • Local replication with TimeFinder SnapVX • Embedded NAS (eNAS) • Embedded Unisphere for PowerMax (eManagement)

Claims	Exemplary Evidence of Infringement

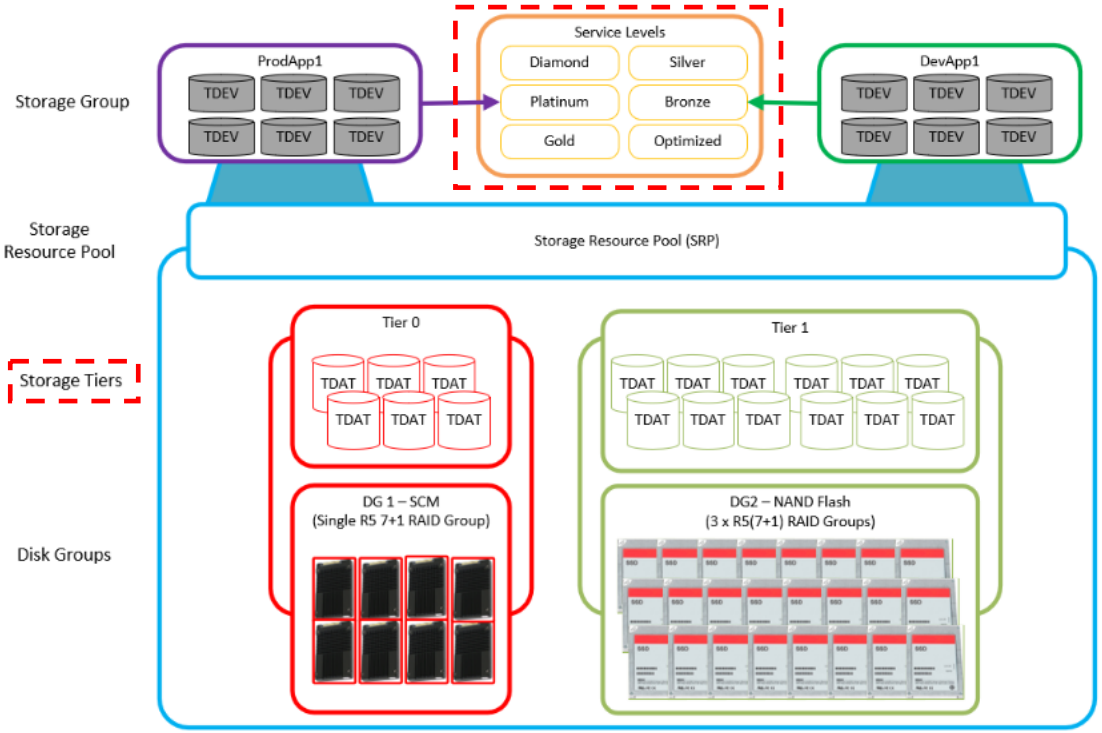

Figure 4. Typical components found with a PowerMax SRP with example of disk-group RAID-protection schemes

Next-Generation PowerMax Family Overview

Claims	Exemplary Evidence of Infringement
	PowerMax Storage Resource Pools overview In PowerMax, all physical storage capacity is combined into Storage Resource Pools (SRPs). At the lowest levels, SRPs consist of Disk Groups which contain of a collection of hard drives sharing the same technology and performance characteristics. The hard drives in each disk group are split into individual back-end data device segments called TDATs. The TDATs are placed into an associated Storage Tier. An SRP is the collection of the total capacity of all its Storage Tiers – regardless of the underlying disk technology which the storage tiers are associated with. This physical capacity stored within an SRP is referred to its usable capacity (TBu). This usable capacity is accessed by hosts using thinly provisioned front-end storage devices called TDEVs. TDEVs are virtual representation of the SRP physical capacity which also considers overprovisioning and data reduction efficiencies. For example, an array with a single SRP which has 26 TBu, could be provisioned for 78 TB of host facing TDEV capacity when a data reduction ratio of 3:1 is applied. This 78 TB of virtualized host facing TDEV capacity is referred to be the effective capacity (TBe) of the SRP. When a PowerMax is sized, both the usable capacity and effective capacity are considered. The total usable capacity (TBu) is the primary driver for sizing hard-drive-layout configurations. The effective capacity (TBe) is a primary driver when sizing PowerMax cache. Host provisioned TDEVs to are placed into a storage group and assigned a Service Level. When a host writes application data to its provisioned TDEVs, this data is distributed across all the storage tiers within the SRP. Which storage tier the data is placed on within the SRP is governed by the Automated Data Placement (ADP) utility. ADP uses the PowerMax internal machine learning engine to employ predictive analytics and pattern recognition algorithms to place the data at the optimal physical location to ensure that the response time requirements for the assigned service level are met. Dell EMC PowerMax: Family Overview

Claims	Exemplary Evidence of Infringement

An SRDF device is a logical device paired with another logical device that resides in a second array. The arrays are connected by SRDF links. R1 devices are the member of the device pair at the primary (production) site. R1 devices are generally read/write accessible to the host. R2 devices are the members of the device pair at the secondary (remote) site. During normal operations, host I/O writes to the R1 device are mirrored over the SRDF links to the R2 device.

Figure 1. Traditional SRDF device pair states

Traditionally, data on R2 devices are not available to the host while the SRDF relationship is active. In SRDF synchronous mode, an R2 device is typically in read-only mode (write disabled) that allows a remote host to read from the R2 devices. In a typical open systems host environment, the production host has read/write access to the R1 device. A host connected to the R2 device has read-only access to the R2 device. To access the R2 device of a traditional synchronous relationship, a manual failover or swap operation must be performed to write enable the R2 site to accept host writes.

Dell PowerMax and VMAX All Flash: SRDF/Metro Overview and Best Practices

Claims	Exemplary Evidence of Infringement
13 [b] [iii] concurrently read data during the read operations from the number of the storage elements associated with the performance index and not currently being used by the write operations.	The Accused Products concurrently read data during the read operations from the number of the storage elements associated with the performance index and not currently being used by the write operations. For example, the Accused Products combine "all physical storage capacity into Storage Resource Pools (SRPs)" which "consist of Disk Groups which contain [] a collection of hard drives." For example, the "hard drives in each disk group are split into back-end data segments called TDATs" and "TDATs are placed into an associated Storage Tier." For example, the "useable capacity of" an "SRP" is "accessed by hosts using thinly-provisioned front-end devices called TDEVs," where "TDEVs are [a] virtual representation of the SRP physical capacity" and "are placed into a storage group and assigned a Service Level." For example, the Accused Products provide "PowerMax Data Services" including "[r]emote replication" with the "Symmetrix Remote Data Facility (SRDF)" and "[h]igh availability with SDRF/Metro." For example, a "SRDF device is a logical device paired with another logical device that resides in a second array" and "[d]uring normal operations, host I/O writes to the R1 device are mirrored over the SRDF links to the R2 device." For example, in "SRDF synchronous mode, an R2 device is typically in read-only mode (write disabled) that allows a remote host to read from the R2 devices" and a "host connected to the R2 device has read-only access to the R2 device." For example, the Accused Products include a "write_disable . . . operation [that] sets source . . . devices as write disabled to their local hosts." For example, "[u]sers may set the required service level objective (SLO) independently on both source and target SRDF/Metro paired arrays." *See, e.g.*: • High Availability with SRDF/Metro • Local replication with TimeFinder SnapVX • Embedded NAS (eNAS) • Embedded Unisphere for PowerMax (eManagement) • Remote replication with SRDF • High Availability with SRDF/Metro • Local replication with TimeFinder SnapVX • Embedded NAS (eNAS) • Embedded Unisphere for PowerMax (eManagement)

Claims	Exemplary Evidence of Infringement

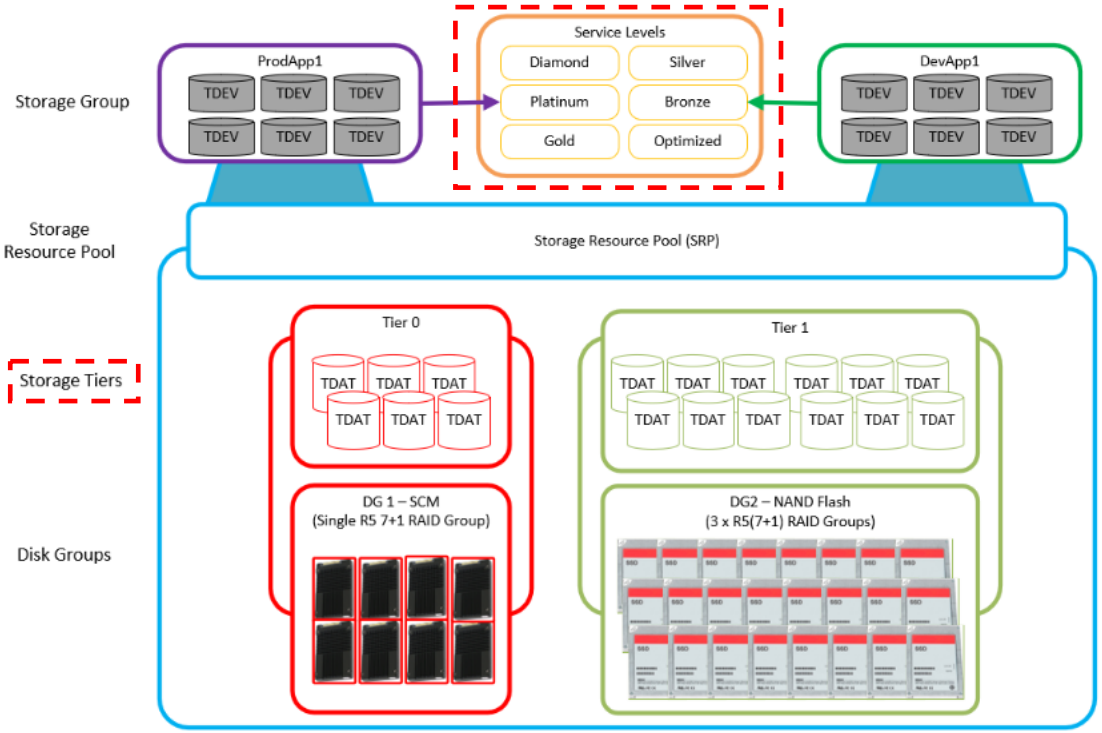

Figure 4. Typical components found with a PowerMax SRP with example of disk-group RAID-protection schemes

Next-Generation PowerMax Family Overview

Claims	Exemplary Evidence of Infringement
	PowerMax Storage Resource Pools overview In PowerMax, all physical storage capacity is combined into Storage Resource Pools (SRPs). At the lowest levels, SRPs consist of Disk Groups which contain of a collection of hard drives sharing the same technology and performance characteristics. The hard drives in each disk group are split into individual back-end data device segments called TDATs. The TDATs are placed into an associated Storage Tier. An SRP is the collection of the total capacity of all its Storage Tiers – regardless of the underlying disk technology which the storage tiers are associated with. This physical capacity stored within an SRP is referred to its usable capacity (TBu). This usable capacity is accessed by hosts using thinly provisioned front-end storage devices called TDEVs. TDEVs are virtual representation of the SRP physical capacity which also considers overprovisioning and data reduction efficiencies. For example, an array with a single SRP which has 26 TBu, could be provisioned for 78 TB of host facing TDEV capacity when a data reduction ratio of 3:1 is applied. This 78 TB of virtualized host facing TDEV capacity is referred to be the effective capacity (TBe) of the SRP. When a PowerMax is sized, both the usable capacity and effective capacity are considered. The total usable capacity (TBu) is the primary driver for sizing hard-drive-layout configurations. The effective capacity (TBe) is a primary driver when sizing PowerMax cache. Host provisioned TDEVs to are placed into a storage group and assigned a Service Level. When a host writes application data to its provisioned TDEVs, this data is distributed across all the storage tiers within the SRP. Which storage tier the data is placed on within the SRP is governed by the Automated Data Placement (ADP) utility. ADP uses the PowerMax internal machine learning engine to employ predictive analytics and pattern recognition algorithms to place the data at the optimal physical location to ensure that the response time requirements for the assigned service level are met. Dell EMC PowerMax: Family Overview

Claims	Exemplary Evidence of Infringement

An SRDF device is a logical device paired with another logical device that resides in a second array. The arrays are connected by SRDF links. R1 devices are the member of the device pair at the primary (production) site. R1 devices are generally read/write accessible to the host. R2 devices are the members of the device pair at the secondary (remote) site. During normal operations, host I/O writes to the R1 device are mirrored over the SRDF links to the R2 device.

Figure 1. Traditional SRDF device pair states

Traditionally, data on R2 devices are not available to the host while the SRDF relationship is active. In SRDF synchronous mode, an R2 device is typically in read-only mode (write disabled) that allows a remote host to read from the R2 devices. In a typical open systems host environment, the production host has read/write access to the R1 device. A host connected to the R2 device has read-only access to the R2 device. To access the R2 device of a traditional synchronous relationship, a manual failover or swap operation must be performed to write enable the R2 site to accept host writes.

<table>
<tr><th>Claims</th><th>Exemplary Evidence of Infringement</th></tr>
<tr><td></td><td>
Performance statistic exchange begins once the SRDF/Metro Active mode and ActiveActive or ActiveBias pair state have been achieved. Each side then incorporates the FAST statistics from the other side to ensure each side represents the workload as a whole (R1+R2 workload). Users may set the required service level objective (SLO) independently on both source and target SRDF/Metro paired arrays. There are currently no restrictions in this area as FAST data movement is transparent from SRDF/Metro.
<table>
<tr><th>State</th><th>Description</th></tr>
<tr><td>SyncInProg</td><td>Synchronization is currently in progress between the R1 and the R2 devices.
• Percent Done: not 100% (The background copy is not complete)
• The copy direction could be R1→R2 or R1←R2.
• RDF Device Link Status: Ready (RW)
• MetroR1, host accessible
○ Device RDF State: Ready (RW)
○ Device SA Status:
▪ If mapped, Ready (RW)
▪ If unmapped, N/A
○ Device Status: Ready (RW)
• The MetroR2, not host accessible
○ Device RDF State: Write Disabled (WD)
○ Device SA Status:
▪ If mapped, Ready (RW)
▪ If unmapped, N/A
○ Device Status: Not Ready (NR)</td></tr>
</table>
Dell PowerMax and VMAX All Flash: SRDF/Metro Overview and Best Practices
</td></tr>
</table>

Claims	Exemplary Evidence of Infringement

Figure 11 Update SRDF device track tables

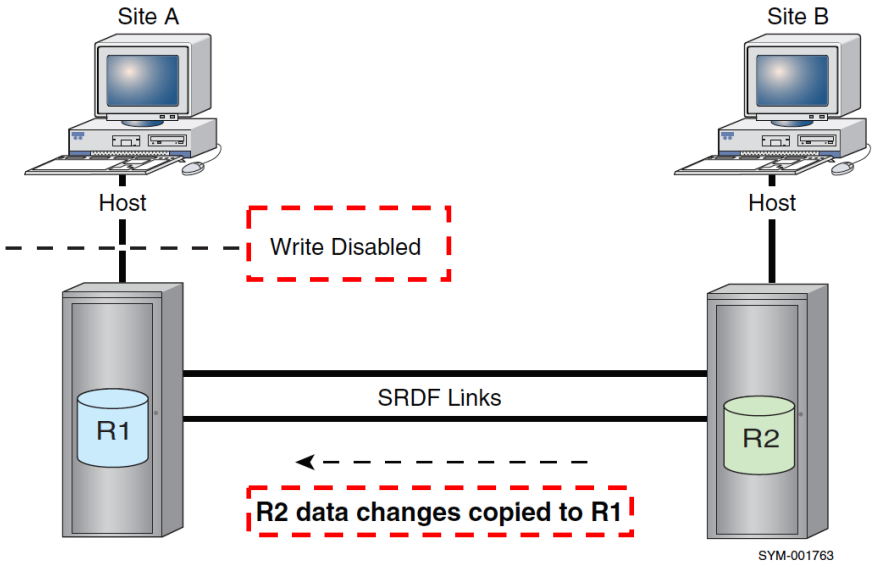

Note

When you issue the `symrdf` command, device external locks are set on all SRDF devices you are about to control. See Device external locks on page 56 and Table 5 on page 37.

Write enable R1

The read/write enable R1 operation makes the source (R1) devices accessible to their local hosts.

Claims	Exemplary Evidence of Infringement

Syntax

Use `rw_enable r1` for a device group, composite group, or device file:

```
symrdf -g DgName rw_enable r1
symrdf -cg CgName rw_enable r1
symrdf -f[ile] FileName rw_enable r1
```

Examples

To enable all the source (R1) mirrors in all the SRDF pairs in device group `prod`:

```
symrdf -g prod rw_enable r1
```

Write disable R1

The `write_disable R1` operation sets the source (R1) devices as write disabled to their local hosts.

Syntax

Use `write_disable r1` for a device group, composite group, storage group, or device file:

```
symrdf -g DgName write_disable r1
symrdf -cg CgName write_disable r1
symrdf -sg SgName write_disable r1
symrdf -f[ile] FileName write_disable r1
```

Examples

To write disable all the source (R1) mirrors in the SRDF pairs in device group `prod`:

```
symrdf -g prod write_disable r1
```

Dell EMC Solutions Enabler SRDF Family Version 9.0 CLI User Guide